

188083

S1-63

7P.

NC473657

## **Building Intelligent Systems: Artificial Intelligence Research at NASA Ames Research Center**

**P. Friedland and H. Lum**  
NASA Ames Research Center  
Moffett Field, CA 94035

### **1. Introduction**

The goal of building autonomous, intelligent systems is the major focus of both basic and applied research in artificial intelligence within NASA. This paper discusses the components that make up that goal and describes ongoing work at the NASA Ames Research Center to achieve the goal.

NASA provides a unique environment for fostering the development of intelligent computational systems; the test domains, especially Space Station systems, represent both a well-defined need, and a marvelous series of increasingly more difficult problems for testing research ideas. One can see a clear progression of systems through research settings (both within and external to NASA) to space station testbeds to systems which actually fly on space station.

### **2. The Long-Range Goal**

As a springboard for discussion, let us create a view of a "truly" autonomous space station intelligent system, responsible for a major portion of space station control (the exact system is unimportant). We will build a view of all of the functions this system should have, which of those functions we can achieve (nearly completely) today, which we can easily see happening in the next few years with some engineering-oriented applied research, and which will be doable only with substantial basic research (over at least the next five years).

Our intelligent system will have full responsibility for a major functional component of space station; examples include power, communications, thermal management, and environmental control. It will be responsible for nominal control, acute and chronic failure discovery, diagnosis, and correction, and communications/cooperation with both interested humans (astronauts, ground controllers, scientists) and other intelligent computational systems. The following specific abilities are needed to be able to completely satisfy those responsibilities:

- Scheduling of System Resources to Meet Utilization Requests--the ability to analyze tasks that involve the space station component being controlled in order to set up targets for resource utilization.
- Real-Time Schedule Execution and Monitoring--the ability to translate task requests into executable system commands and to understand if the tasks have been adequately performed.
- Dynamic Schedule Modification--the ability to change the resource utilization schedule to reflect both internal (poorly performed or understood tasks) and external constraints (new tasks added, conflicts with other space station systems, etc.).
- Heuristic/Experiential Failure Detection and Diagnosis--the ability to utilize "shallow" knowledge based mainly upon prior (human or machine) experience with the space station system in order to notice and diagnosis problems with the system.

- **Causal/Model-Based (Both Qualitative and Quantitative) Failure Diagnosis**--the ability to utilize deeper, "first-principle" knowledge to diagnose system problems.
- **Planning of Failure Corrections**--the ability to determine a course of actions to repair a diagnosed failure.
- **Realtime Failure Correction and Monitoring**--the ability to translate a correction plan into actions and understand if the problem has, indeed, been fixed.
- **Long-Term Trend Analysis**--the ability to understand slow-to-develop trends either to prevent incipient failure or to adjust nominal control over a long period of time.
- **Explanation of Actions**--the ability to explain, with clarity and brevity, all system actions, from resource utilization schedules to failure diagnosis to failure correction actions. Clarity and brevity imply different explanations for different classes of humans who wish to interact with the system.
- **Cooperation with Other Intelligent Systems**--the ability to work in concert with other intelligent entities on space station, both human and computational.

Four other necessary abilities cross many functional lines for our truly autonomous system. These are:

- **Reasoning Under Uncertainty**--the ability to make sensible judgments and carry out reasonable actions when world knowledge is imprecise or incomplete, heuristics or models have built-in uncertainty, or actions have uncertain effects.
- **Learning**--the ability to alter and improve all functionalities as conditions change and knowledge is added over time. Starting views of how systems operate in space may be wrong. Fault correction actions may alter system configurations. An autonomous system cannot remain static or performance will at best not improve or (more likely) at worst degrade to completely unacceptable.
- **Common Sense Reasoning**--the ability to occasionally go beyond specific domain knowledge into broad areas of human experience. This includes the ability to bypass established reasoning mechanisms when unexpected events render them clearly useless.
- **Self-Understanding**--the ability to understand:
  - When and how to utilize different functional abilities such as heuristic as opposed to model-based diagnostic reasoning.
  - When to act independently as opposed to asking for human or other computational assistance.
  - How to prioritize actions--when are things critical and when can they wait.
  - When a problem is beyond the system's range of understanding.
  - When new knowledge is worth saving either directly or as part of a new plan or piece of a model.

Finally, to build such intelligent, autonomous systems, three pieces of artificial intelligence methodology become critical. These are:

- **Better Knowledge Acquisition**--great improvement in our abilities to acquire, both initially and dynamically via sensor and other input, the heuristic and causal knowledge our system requires. In other words, tools and techniques for constructing the knowledge base.

- Managing of Large Knowledge Bases--how to manage, optimize, check for internal consistency and completeness, etc., knowledge bases at least a few orders of magnitude larger than those used by current large systems.
- Validation of Intelligent Systems--both practical and philosophical understanding of what it means to validate autonomous, computational systems. For space station systems, a mutual education and satisfying process with Space Station personnel and the aerospace industry will be part of the task.

### 3. Where We Are

The state of the art is marked by the three orthogonal areas of kinds of knowledge we can represent, tools for knowledge base construction and manipulation, and kinds of problem solving. Using one of the three representation paradigms of rules, frames, and logic, we have engineering solutions for storing a wide variety of heuristic/procedural knowledge and factual/declarative knowledge. We have commercial tools (KEE, S.I., KnowledgeCraft and ART are the outstanding examples) which reduce the knowledge base construction time by at least two orders of magnitude over raw LISP (or FORTRAN, PROLOG, etc. for that matter) and which provide a selection of inference methodologies. We can routinely carry out the entire knowledge-based system building task for structured selection problems using experiential knowledge, component configuration problems, fairly complex scheduling problems, simple planning problems, and "intelligent front-ends" for abstruse modeling and database computer programs.

In addition, we can represent and utilize a very limited amount of probabilistic or conditional knowledge about data and knowledge. Our systems can explain their chain of reasoning in simplistic textual and graphical forms. We can also construct quite sophisticated, personalized interfaces to knowledge-based systems. Finally, we can utilize precisely the same knowledge base for distinct purposes, including diagnosis, simulation, and training.

### 4. What We Can Almost Do

In several other areas, basic research is beginning to make the transition into engineering utilization. Specifically:

- Causal models are beginning to be used on significant problems, augmenting and/or replacing experiential heuristics for diagnosis and fault correction.
- Systems are moving from "leisurely" offline applications into realtime control systems (where realtime currently means on the order of seconds for evaluation and response time).
- Systems that dynamically reconfigure plans to reflect changing conditions during plan execution are beginning to appear in real applications.
- The blackboard framework for cooperative utilization of different sources of knowledge is becoming part of the inference "toolkit." This represents a "micro" view of distributed control among cooperating knowledge-based systems in most current applications.
- Machine learning, still in very simple forms such as explanation based generalization, or learning by example, is starting to make an impact on fielded knowledge-based systems.

## 5. What We Need to Do

Clearly there is an enormous amount of work to be done to take us from where we are (or soon will be) in the technology of building autonomous systems to a point where the long-range goal is satisfied. Just as clearly, we at NASA Ames cannot do it all. However, the problem areas discussed in this document are relevant to most current work in artificial intelligence research, so we have a wide community of fellow researchers to draw upon. The following are the scientific and engineering research areas in which we at NASA Ames believe our involvement through internal work or support of external work makes sense over the next several years:

- Reasoning Under Uncertainty--as discussed above, particularly focusing on integration into practical knowledge acquisition and representation frameworks and demonstrations of utility.
- Machine Learning--emphasizing automatic knowledge base expansion and correction as well as learning by discovery (carrying out sensor-based "experiments, etc.).
- Causal Modeling and Simulation--particularly on methodologies for integrating these methods with both heuristic-based problem solving at one end and mathematical model based simulation at the other.
- Next Generation Tools for Knowledge Acquisition, Representation, and Manipulation--developing better, faster, more versatile tools to "routinize" much of the knowledge engineering process. Of particular importance is getting past the knowledge acquisition bottleneck and removing a human knowledge engineer from the loop as much as possible.
- Explanation and Interface Technology--making progress in communicating between knowledge-based systems and human users. We wish to add perspicuity and perspicacity of explanation, and make it much more easy to customize interfaces to individual preferences.
- Constructing and Utilizing Large Knowledge Bases--fundamental experiments in treating very large collections of knowledge.
- Acquisition of Design Knowledge and Data for Complex Systems--ensuring that the vast amount of information used and discovered during the design process not be lost for future knowledge-based systems that need to reason about the artifact that was constructed.
- Advanced Methods for Plan Construction, Monitoring, and Modification--integration of current planning methods to ensure usability and flexibility within practical environments
- Hierarchical Control of Multiple Knowledge-Based Systems--experiments in control of multiple systems by hierarchical levels of increasingly more general knowledge-based systems.
- Distributed Cooperation among Multiple Knowledge-Based Systems--the alternate view to the previous topic: coordinated control by cooperative information sharing through a common database or "blackboard."
- Validation Methodologies for Knowledge-Based Systems--both low-level issues of software verification (eliminating redundancies, inconsistencies, etc.) and more fundamental issues of ensuring desired functionality.

Note that certain topics relevant to the Long-Range Goal, in particular common sense reasoning and self-understanding, are not specifically covered in the above research topics. That is because, in the author's view, meaningful work in those areas awaits fundamental

progress in other areas like causal modeling and machine learning. There are several very long-term research projects already underway in the "missing" research areas (Lenat's CYC project at MCC is a good example) and we will keep careful track of those projects as part of our overall understanding of autonomous systems.

## **6. Getting It Done**

The previous section has established ambitious goals for our research in intelligent systems. Over the next several years, we will be attempting to build a highly regarded basic research laboratory within the Artificial Intelligence Research and Applications Branch at Ames. As relevant personnel join the program, research projects in many of the above areas will commence. Indeed, as discussed below, several projects already exist or will begin in fiscal 1987. However, for the next year or so, our greatest leverage will come from sponsoring a modest number of extremely high quality external research projects. The external work, besides serving programmatic research needs, will also, we believe, act to educate Ames personnel through collaborative discussions, produce some short-term technology demonstrations, and provide a stream of motivated students, a significant number of whom we hope can be convinced to join NASA upon receipt of their advanced degrees. In addition, we will provide attractive mechanisms for senior researchers to carry out portions of their NASA-sponsored work at the Ames Research Center through sabbaticals and summer visitations.

The final two sections of this document will briefly describe current work in intelligent systems in progress and sponsored by the Artificial Intelligence Research Branch at NASA Ames Research Center.

### **6.1. Internal Research in Fiscal 1987**

#### **Reasoning Under Uncertainty**

- Peter Cheeseman, John Stutz, Mary Duffy, et. al. will continue their work in probabilistic reasoning with test applications in planning and learning. The coming year will see a focus on NASA problem domains and comparisons with other schools of uncertain reasoning. Long term goals revolve around developing engineering methodologies for routinely including uncertain reasoning in relevant applications across many different forms of problem solving.

#### **Machine Learning**

- Michael Sims and Peter Friedland will initiate a project on learning by discovery. The starting point will be Sims' work in mathematical discovery applied to the engineered systems such as those on Space Station. The eventual goal is to achieve self-improving knowledge-based control and analysis systems.

#### **Causal Modeling and Simulation**

- William Erickson, Mary Schwartz, and Peter Friedland will use the space station thermal system as a testbed for work in integrating causal modeling with experiential heuristics and mathematical models. A short term goal is to demonstrate practical accomplishments as part of the 1988 Systems Autonomy Demonstration Program. A somewhat longer-term goal is to formalize the techniques for determining when each of the different forms of reasoning will be most useful for difficult problems.

#### **Knowledge from Design Through Operations**

- Cecelia Sivard and Lilly Spirkovska will continue their work on capturing design knowledge on the solar photometer system. A short-term goal is a proof of utility of the methodology, by showing how knowledge acquired during the design process

can be effectively utilized during operations and maintenance of a complex device. A long-term goal is building increasingly more sophisticated design tools that automatically construct knowledge bases useful during the entire design through operations process.

#### **Advanced Planning Work**

- Peter Cheeseman, et. al. will continue their work on dynamic scheduling and rescheduling. Peter Friedland will collaborate to see if skeletal planning methodologies can be integrated with the work. The eventual goal is to show realtime utility of the methodologies.

#### **Validation Methodologies**

- Peter Friedland and Carla Wong will begin a project in the practical aspects of knowledge-based system validation. This will include intense collaboration with Space Station and the space industry to prevent bottlenecks in the goal of a clear path to operational systems on space station. First results will be seen in the 1988 Systems Autonomy demonstration. A longer-term goal will be NASA-wide accepted methodologies for knowledge-based system validation.

### **6.2. External Research in Fiscal 1987**

#### **Reasoning Under Uncertainty**

- Lotfi Zadeh, UC-Berkeley-- Dr. Zadeh is the world leader in the branch of uncertain reasoning known as fuzzy logic. During the coming year, we will begin the process of integrating his work with Ames-internal efforts.
- Don Heckerman and Eric Horvitz, Knowledge Systems Laboratory, Stanford University--Heckerman and Horvitz are students who have become leaders in both theoretical and practical aspects of the uncertain reasoning work that originated in the MYCIN project. During the next year they will begin to determine the relevance of those methodologies to NASA domains.

#### **Machine Learning**

- Tom Mitchell and Jaime Carbonell, Carnegie-Mellon University--Mitchell and Carbonell, are among the world's leading authorities on research in machine learning. The work will emphasize explanation based generalization and learning by example and include both theoretical aspects and practical applications to NASA domains. In addition, both Mitchell and Carbonell are likely to spend significant amounts of time at the Ames Research Center helping build our internal strengths in machine learning.

#### **Advanced Planning Work**

- David Atkinson, JPL--The AI research group at JPL will be conducting Ames-sponsored research in the integration of sensor-based planning, plan monitoring, and plan modification. This is the first effort in what we hope will be a long-term research collaboration with the JPL group.

#### **Hierarchical Control of Multiple Knowledge-Based Systems**

- Ron Larsen, University of Maryland--We will continue to support Dr. Larsen's work on developing computational structures for hierarchical control of cooperating knowledge-based systems. His work is attempting to integrate traditional decision theory with heuristic and model-based control methods.

- Richard Volz, University of Michigan--Professor Volz is carrying out a variety of projects in sensor-based learning, manipulation, and scene understanding related to coordinated control of robotic systems. This research is being sponsored in cooperation with the JPL-led Telerobotics Program.

#### **Distributed Cooperation among Multiple Knowledge-Based Systems**

- Ed Feigenbaum and Bruce Buchanan, Knowledge Systems Laboratory, Stanford University--The Knowledge Systems Laboratory is conducting a significant new effort centering on methods for cooperative control among distributed expert systems. This work includes both long-term research projects resulting in Ph.D. theses and short term development of blackboard-based techniques for cooperative, distributed control.
- Tom Sheridan, MIT--We have been supporting Professor Sheridan's work in various methodologies for distributed and cooperative control in robotic systems. This work will continue as a joint project with the JPL-led Telerobotics Program.